Networking Basics

EC512

Spring 2015

Protocols

- Protocols are required in order to allow information to be extracted from the stream of bits flowing from one point in a data communication link to another.

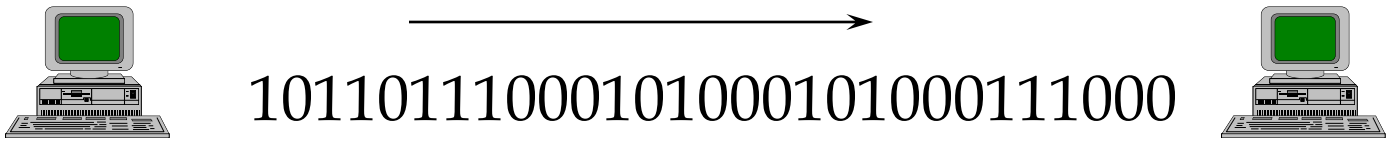

Encapsulation

- Protocols encapsulate data inside a ***frame***.
- A *frame* consists of a *header*, *data*, and an optional *trailer*.

Encapsulation - contd.

- The header contains a variety of information which may include:
  - Source and destination addresses.
  - Size of the data.
  - Protocol used by the data.
  - Error checking such as a checksum/CRC
- A trailer is sometimes used to hold error checking information if it is not included in the header.

Maximum Transmission Unit - MTU

- Every protocol has an associated MTU.
- This is the maximum size of the data that may be encapsulated in the frame.
- The MTU size is set by a variety of factors.
- The MTU for Ethernet is typically 1500 bytes. Serial connections normally use smaller MTUs.

Layering

- It is not practical to have just one protocol that does everything.
- Different protocols are used at different layers.
- The OSI model has seven layers.
- The TCP/IP protocol suite has four layers.

The TCP/IP Layers

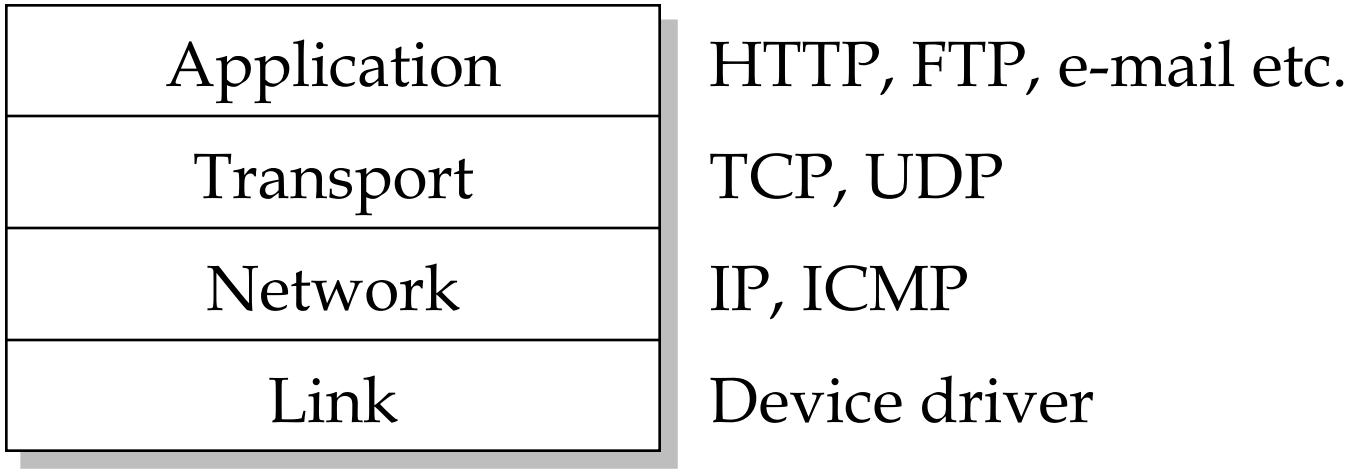

How Hosts Communicate

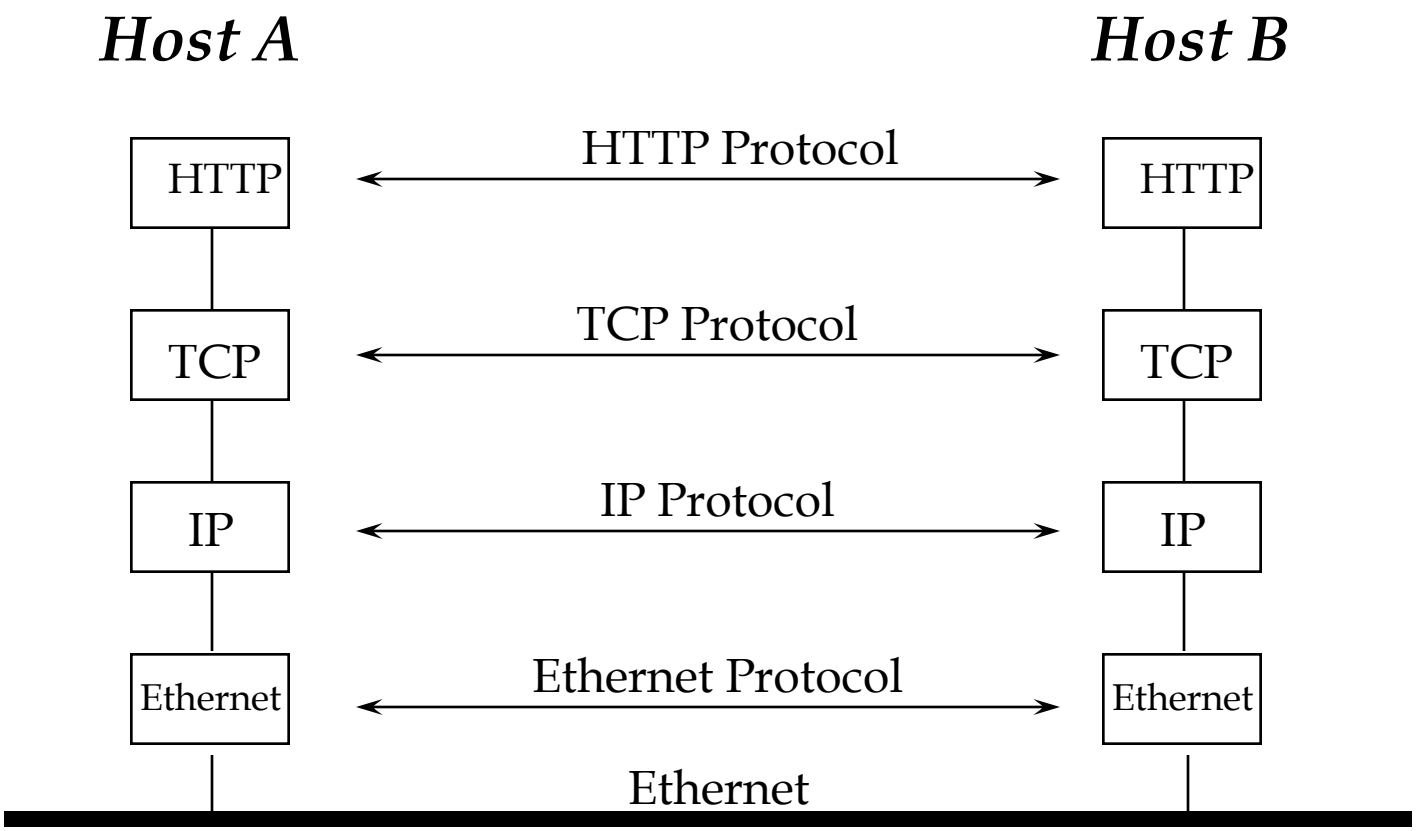

Encapsulation - contd.

- A complete frame at one protocol layer becomes the *data* for the next lower layer.

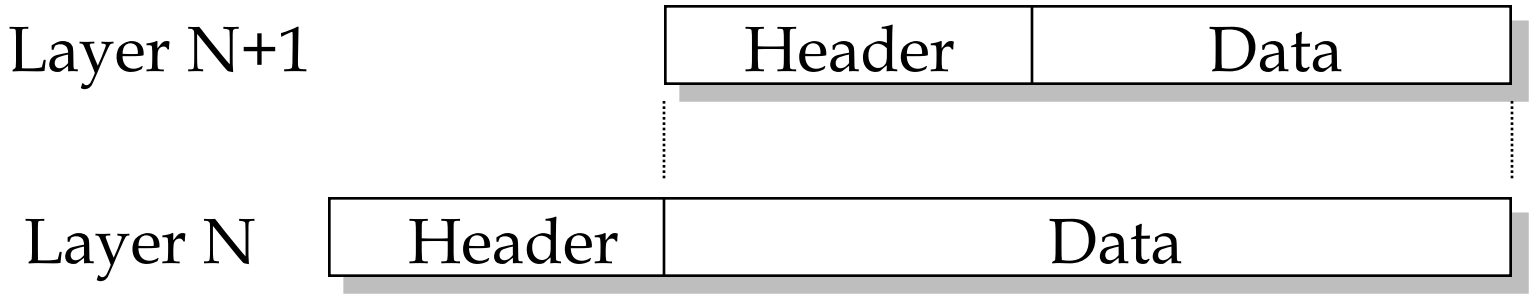

Peer Layer Communication

- Layer N has an *interface* to layer N-1 and to layer N+1 on each host.
- Layer N on one host communicates with the corresponding layer N on the other host.
- Different protocols can be used at the same layer, but they must be the same at ***both*** ends for a specific application.

Peer Layer Communication - contd.

- The single exception to this rule for the TCP/IP protocol suite is that IP must be used at layer two.
- The IP protocol is the glue that holds the Internet together.
- All hosts. gateways. and routers must be able to communicate using IP.

Demultiplexing

- When an IP datagram is received by a host computer, the datagram is passed to the next higher protocol layer. This process is called ***demultiplexing***.

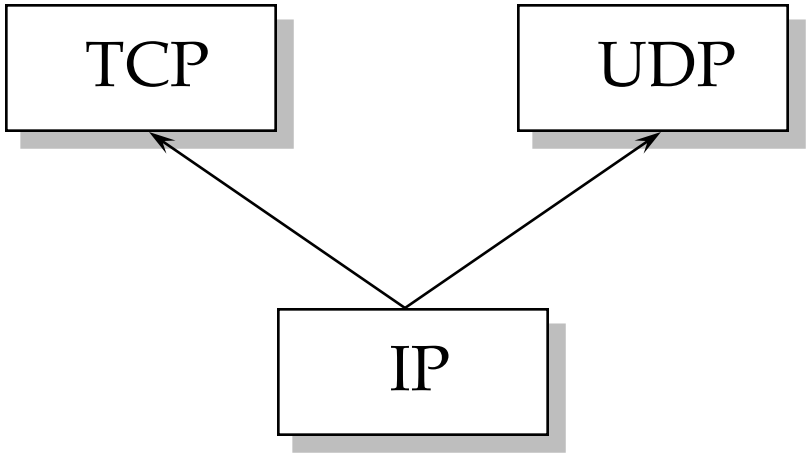

Addressing on the Internet

- An IP address is a 32-bit number. (for IPV4 which dominates the Internet)
- It represents a specific host computer on a specific network. (Actually an interface on a host computer)
- IP addresses are unique over the entire Internet. (Except for hidden and inaccessible local networks, e.g. NAT)
- An IP address consists of a *network* field and a *host* field.

Dotted Decimal Notation

- It is much easier to specify an IP address as four decimal numbers ranging from 0 to 255 rather than a 32-bit binary number.
- However, the break between *netid* and *hostid* is not necessarily on a byte boundary.
- A *Net Mask* is used to specify this division.
- The *Net Mask* determines the number of hosts on the specific subnet)

Ports

- The IP protocol does not have the ability to specify the particular user or program at the destination computer.
- This is similar to delivering a piece of mail to the post office without the box number.
- ***Ports*** provide the ability to send and receive data between applications.
- Port numbers are 16-bits

User Datagram Protocol - UDP

- UDP provides a simple extension to IP to provide ports.
- UDP is used by a variety of Internet applications:
  - Domain name service
  - TFTP file transfer
  - SNMP management protocol etc..
- UDP is *connectionless* and not reliable.

Reserved Ports

- Network services are accessed by the use of reserved ports.
- Reserved ports are port numbers dedicated to specific services.
- For example, HTTP (web browsers) uses TCP port 80. SMTP uses port 25.
- Non-standard ports can be used, but must be advertised as such.
- Ports below 1024 are reserved but ports outside this range may also be reserved on specific operating systems.

Transmission Control Protocol - TCP

- TCP provides:
  - Reliable stream transport.
  - Connection oriented services.
- Data is divided into ***segments.*** Each segment is sent as an IP datagram.
- *Sliding windows* are used to optimize data throughput.
- TCP ports are distinct from UDP ports.

The Domain Name System

- In the early days of networking prior to the Internet (ARPANET days) hosts were assigned unique names. As the number of hosts grew it became more and more difficult to provide unique names.
- In addition, each host needed to update its *host table* periodically in order to locate hosts on the network. This became a monumental task.
- The *domain name system* was developed to provide a solution to this problem.

DNS - contd.

- DNS is a distributed database.
- No single site on the Internet knows all the information.
- DNS is organized hierarchically as a tree of domains.
- Individual organizations can have authority over their own domains.

DNS - contd.

- Countries are assigned two letter top level domain names, e.g. UK for the United Kingdom.
- Domain names in the US are assigned by various authorized registrars (e.g., Network Solutions, Register.com).
- Domain names and IP addresses are not assigned together and a domain may have many different IP network addresses.

DNS - contd.

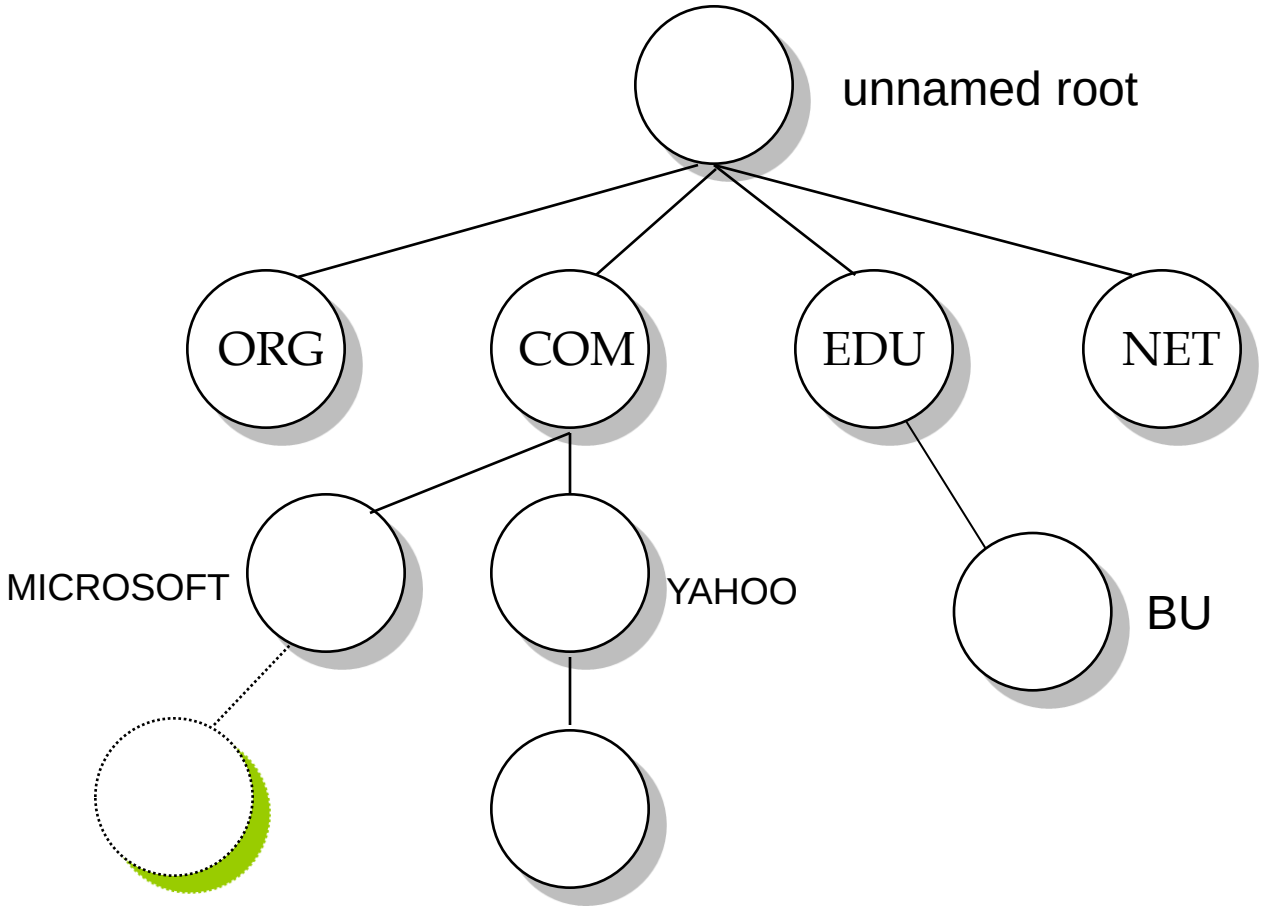

DNS - contd.

- Domains may be broken into sub domains:

  alpha.accounting.xyzcompany.com

  is the host *alpha* in the *accounting* sub domain at *xyzcompany.com*.

DNS - contd.

- Every domain must provide at least two domain name servers that can provide the mappings between domain names and IP addresses.
- One domain name server is the ***primary*** server for that domain. Secondary name servers get their information from the primary. (This is not strictly true for some systems that support *replication* such as Windows Server 2003)
- Updates are done periodically.

DNS - contd.

- A series of ***root*** servers contain the addresses of at least two servers for each domain.
- To *resolve* a domain name to IP address mapping a host sends a query to a domain name server. This is normally one that is provided by an Internet Service Provider (ISP) or a server run by the local domain (for companies).

DNS - contd.

- If the local server has the mapping then it returns the result.
- If the local server does not have the mapping then it queries a root server to get the address of a server for the particular domain requested. It then sends another request to that server for the final mapping.
- This process may continue at yet another level if sub domains are in use.